Xilinx Architects ARM-Based Processor-First, Processor-Centric Device

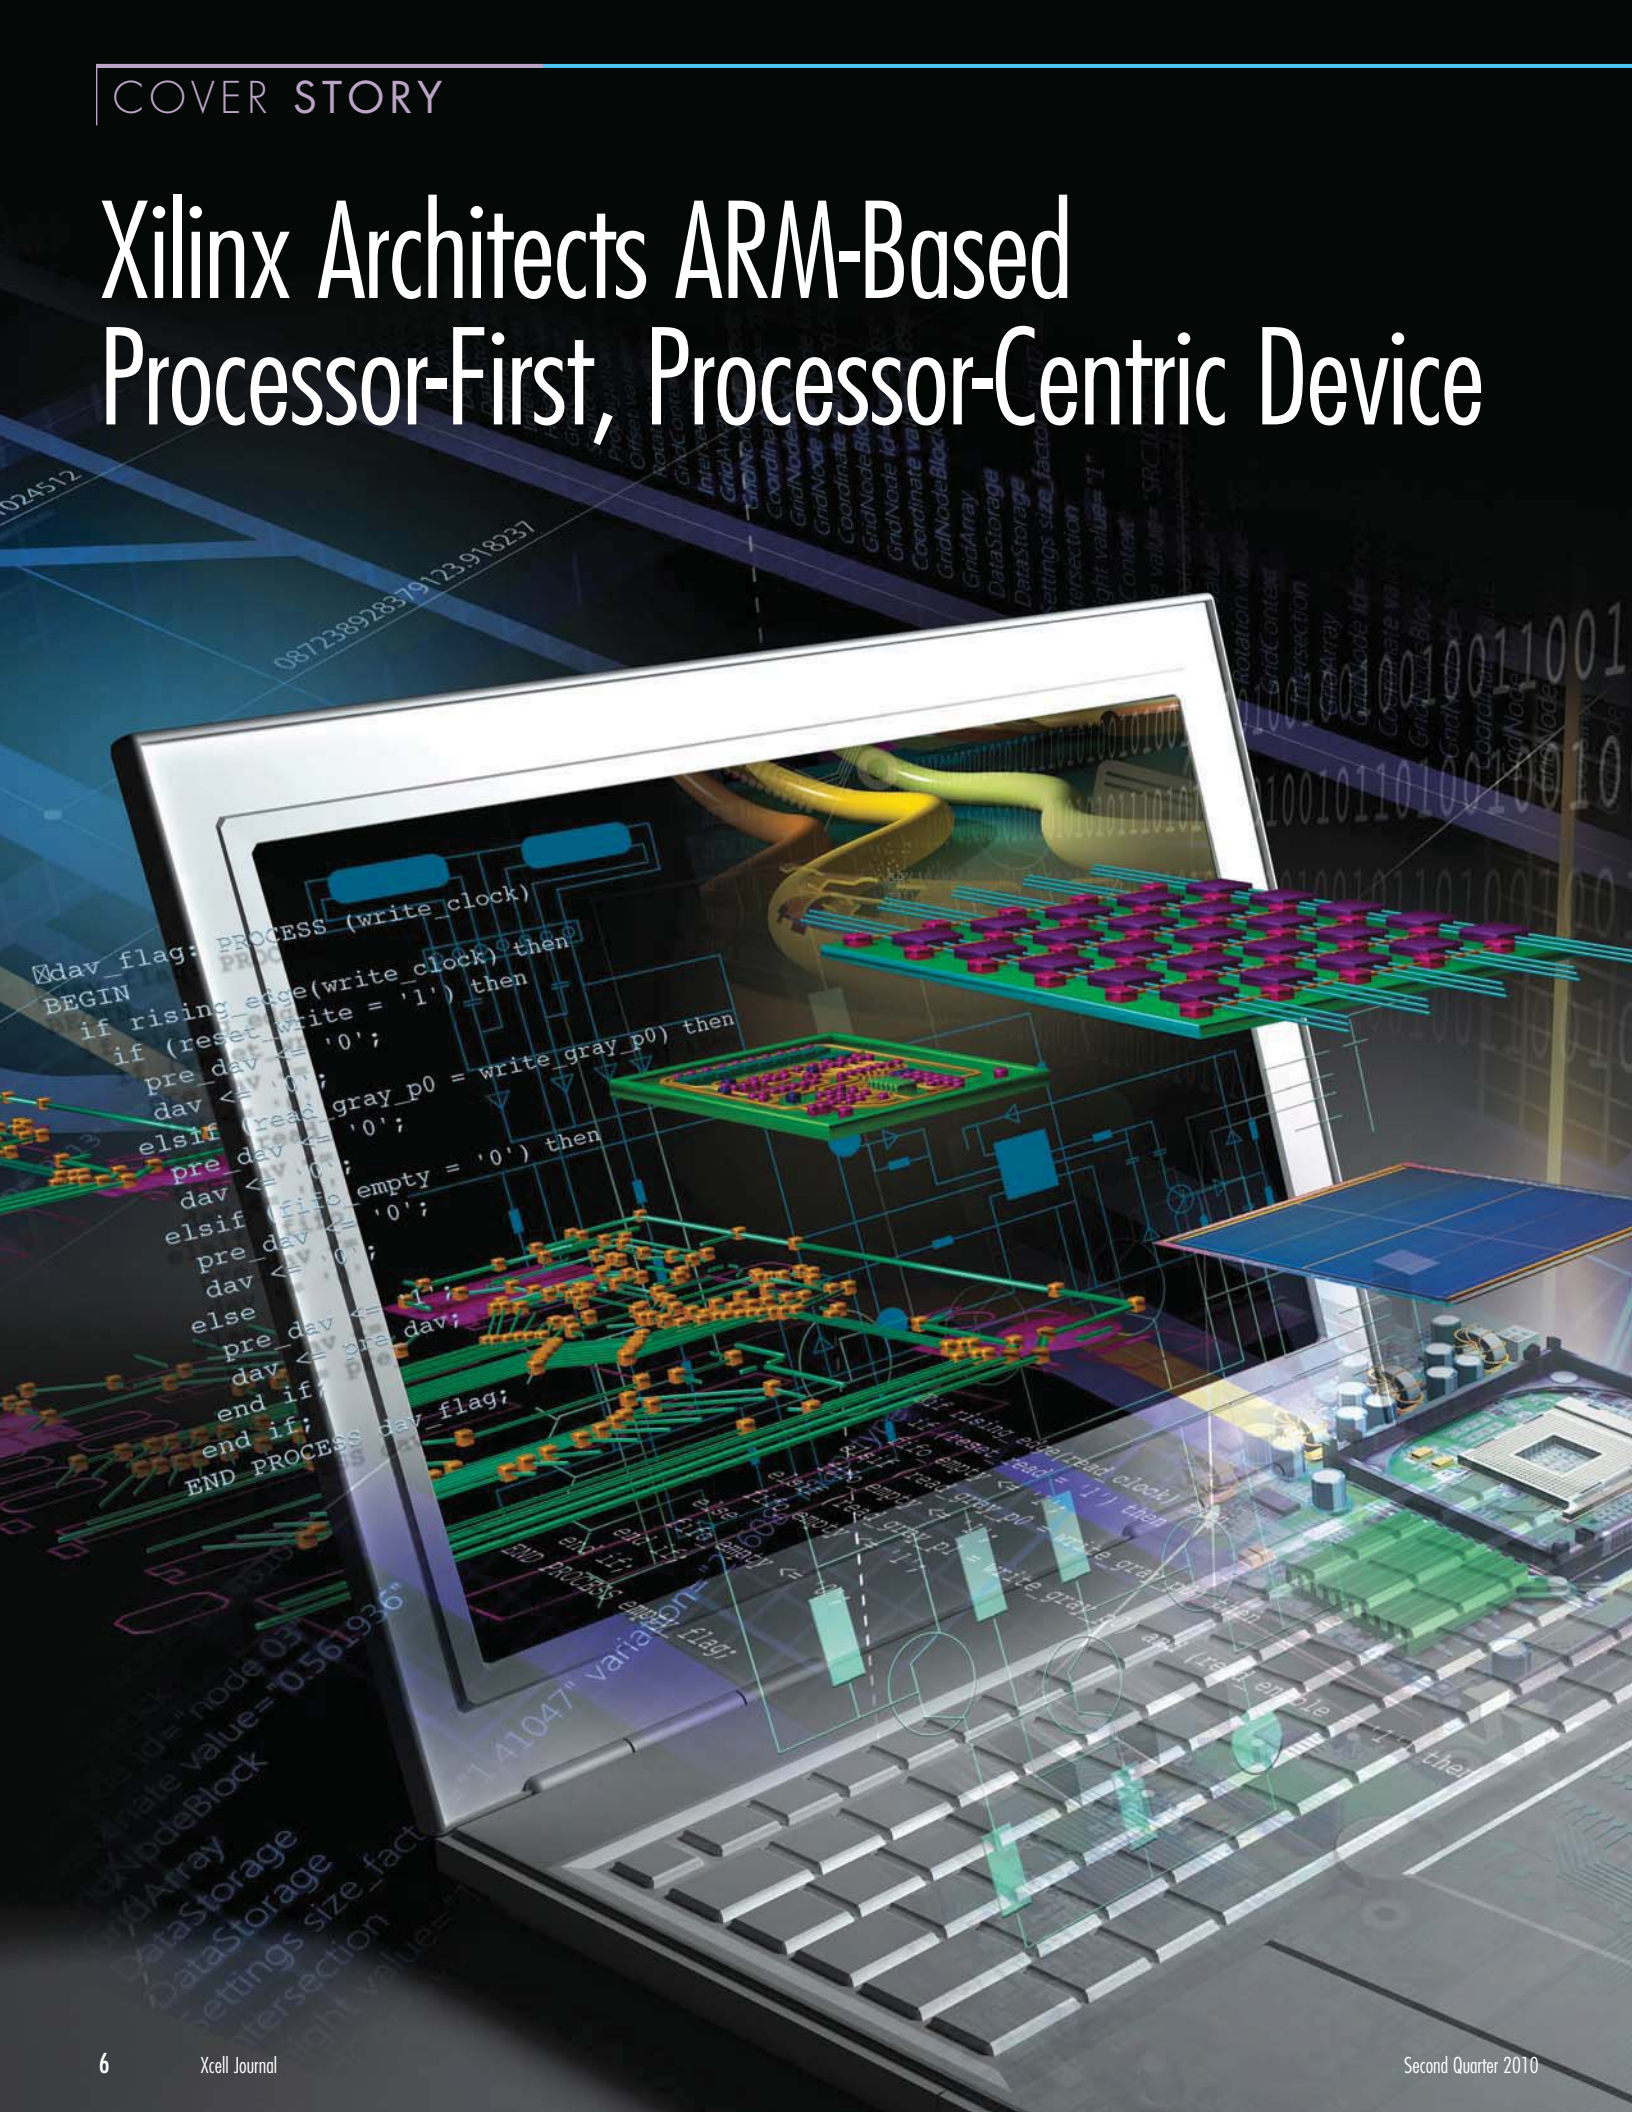

New architecture targets both software and system developers. Processor boots before programmable logic to speed and simplify system development.

by Mike Santarini
Publisher, Xcell Journal
Xilinx, Inc.
mike.santarini@xilinx.com

For well over a decade, FPGA vendors have been searching for ways to add the vast numbers of embedded-software engineers to a user base composed mainly of hardware design engineers. Software developers outnumber their hardware engineer counterparts by up to 10:1 worldwide, so creating a device that both camps can instantly use has obvious business benefits. Now, Xilinx is architecting a new ARM® microprocessor-based device—an Extensible Processing Platform—tailored to the real way software and system developers work. In breaking through the barrier, this novel device promises to take the company into new markets and to new levels of growth.

As FPGA devices and tools have matured over the last decade to incorporate a greater number of embedded processors (DSPs, microcontrollers and microprocessors) in dominantly programmable-logic architectures, a growing number of embedded-system designers have broadened their skill sets beyond middleware and software development and become competent with hardware design languages. This cross-discipline has allowed these happy few designers to begin using FPGAs to create highly optimized, differentiated architectures that have the right mix of hardware and software for prime system performance, functionality and power consumption.

But while the number of engineers with this mixed skill set has grown steadily but slowly over the last 10 years, the vast majority of systems designers still rely on hardware engineers to create custom hardware functions their systems will need to achieve that optimal mix of performance, functionality, power consumption and system cost.

Processor Comes First

With the Xilinx® Extensible Processing Platform, Xilinx is planning to release a first-of-its-kind device in which a 32-bit ARM Cortex™-A9 processing subsystem running at up to 800 MHz is the dominant block (see figure). The processor subsystem will be bootable and programmable out of the box. The balance of the new device will consist of a tightly coupled programmable-logic extension block that will allow designers to partition their hardware and software functions based on system requirements. They can implement functions in the programmable-logic extension block to create their own application-specific, highly optimized systems-on-chips (SoCs).

"We've put a lot of thought and planning into the architecture of the device and have learned a lot of lessons from our previous devices, like our PowerPC™-based Virtex®-II Pro,Virtex-4 and Virtex-5 FXT FPGAs. We've also learned from the missteps of our competitors," said Vin Ratford, senior vice president for worldwide marketing and business development at Xilinx. "Where all those devices took a hardware design-centric view of system design or simply don't have enough processing horsepower, our new Extensible Processing Platform takes a processor-first approach, so software designers can start developing with this product right out of the box. They don't even have to use the extension block if they choose not to."

But many design teams that have a mix of software and hardware designers will embrace the extension block. Xilinx plans to refine the use model over time with the end goal of offering software and system developers an environment that will enable software gurus to program the programmable-logic extension in addition to the processor without the assistance of hardware designers.

Ratford points out that unlike previous architectures in which the FPGA boots before the processor, this new processor-first platform plays perfectly to how developers actually work when building system architectures.

"Electronic systems and software engineers typically develop what they want their system to do in the software first, and then determine what functions they need to accelerate by implementation in hardware," he said. "This allows them to fit their design into the appropriate performance, cost and power footprint that ulti-

mately the application is going to need. When they start a project, they are developing a proof of concept. They are less worried about having it tuned to the specific requirements of a specific customer, and are more concerned about having the maximum amount of flexibility in determining what can be done in hardware or software. Through revisions, they decide what functions will go into hardware and what will go in software, and then take steps to refine both to fit their system requirements. Our device will help them do their job faster and better than before."

Xilinx will deliver the new Extensible Processor Platform in the same low-power, high-performance 28-nanometer process technology that it will use for its next-generation FPGAs (see sidebar).

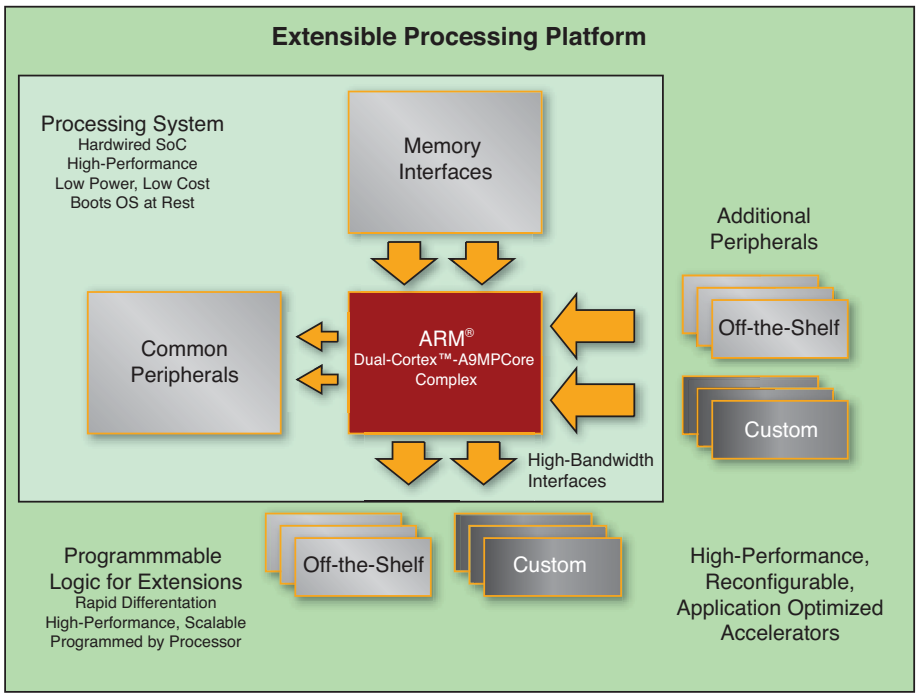

Figure 1 – The Xilinx Extensible Processing Platform pairs an ARM processor with programmable elements.

Why ARM?

Xilinx chose to partner with ARM Ltd. because the company is well-established and has an incredible reputation for the quality of both its processor IP and software. Indeed, the ARM architecture has become the de facto standard choice for designers seeking high-speed, low-power microprocessor cores.

"On all fronts—hardware and software functionality, performance, ecosystem, user familiarity and power—ARM was hands down the best choice for this new architecture," said Ratford. "With power consumption becoming a top-order consideration for not only wireless but also wired applications, adding the lowest-power processor to the device will give users a mind-blowing number of options for making system trade-offs. They can offload functions to the hardware extension block to increase system performance. And they can create systems that can have screaming performance at one moment, then power down to consume milliamps."

A key feature of the new architecture is the interfaces. Xilinx is interconnecting the full, ARM processor-based system and programmable-logic extension block with a high-bandwidth interface connecting the processor-based system, extension block and shared memory. By contrast, a typical system that pairs a discrete MPU-based ASSP chip with an FPGA on the same printed-circuit board typically has just over 100 I/Os connecting them.

Further, in the March release of version 4 of its AMBA® bus' Advanced Extensible Interface (AXI), ARM included an extension of the AXI specification optimized for use on programmable logic. The AXI-4 stream protocol extension serves as a bidirectional crossbar communication switch that will take advantage of the abundant I/Os, enabling engineers using the new Xilinx device to achieve new levels of system interblock throughput and at the same time tap into a multitude of hardware peripheral cores that IP vendors and customers have developed for ARM ASIC and ASSP implementations over the last two decades.

This tightly coupled integration between the two domains, along with the new AXI extension, also means that if design teams find a particular function that isn't running optimally on the processor—or if they have a piece of code they want to speed up—they can create hardware for that function and place it in the programmable-logic extension block using an industry-standard interface.

Familiar Software Programming Model

In creating the new architecture, Xilinx was mindful of the requirements and working preferences of its targeted audience.

Because the new device boots the processor system first, at reset, a software developer can program the processor out of the box, side-by-side with the hardware developer. This has the benefit of shortening development cycles by putting these key functions on parallel tracks.

"In fact, someone could buy the part and just use the processor system," said Keith DeHaven, director of processor marketing at Xilinx. "But the value of the device lies in the user's ability to leverage the programmable logic to customize, optimize and differentiate their products while leveraging the command, control and applications features enabled by the ARM-based processor system."

The real value of the architecture comes as design teams trade off functions between the processor system and programmable-logic extension block. Now the software developer, not the hardware designer alone, will be dictating how the device will operate.

DeHaven said that the processor system has a constant set of peripherals, switches and memory interfaces, providing the software developer with a consistent programming environment. In addition, developers can get started with existing ARM tools and available hardware (summarized in Table 1) to jump-start their efforts.

But of course, the real value of the architecture comes as design teams trade off functions between the processor system and programmable-logic extension block. Now the software developer, not the hardware designer exclusively, will be dictating from the processor view how the device will operate. For example, the processor system may be using data in the extension block to do peripheral functions, or it may delegate control to the block. Developers will likely run hardware/software co-simulation to see if a given function will run in hardware faster, consume less power or reduce cost. Still others may simply want to take software functions that are not likely to change and offload them to the extension block to free up more room in the processor code for other commands.

Once they solidify what functions will go into hardware and software, they can then have their hardware engineers use Xilinx's ISE® Design Suite to implement those functions with an AMBA-AXI standard interface in the programmable-logic extension block. Meanwhile, the developers can continue creating software while the hardware team programs the extension block.

While the processor-first architecture is unique and the use model more representative of how software developers really work, Xilinx plans to make the flow even more intuitive.

Xilinx and its partners are currently developing a comprehensive set of common and standards-based accelerators and peripherals (IP cores, in hardware speak) and related drivers and APIs that will further help software and system developers add functions to their designs. Some of these accelerators and peripherals will be ready at the time of launch, allowing users to focus on creating their own custom IP to meet their system needs and differentiate their products.

The accelerators and peripherals will range in size from small functions that designers can mix and match in the extension block to fully populated extension functions that target specific design domains—connectivity, DSP and processing—and vertical markets such as automotive; industrial, scientific and medical; aerospace and defense; wired and wireless communications, among others.

Vendor	OS
eSol	eT-kernel Multicore Edition
Express Logic	ThreadX
Green Hills	INTEGRITY 10
Kernel.org	Linux 2.6+
Mentor Graphics	Nucleus PLUS RTOS
Microsoft	WinCE
MontaVista	MobiLinux 5.0
QNX	Neutrino RTOS
Symbian	Symbian OS 9+
Wind River	VxWorks 6.6 SMP

Table 1 – ARM has a mature and robust ecosystem for operating systems and OS development tools. Here are some of the OSes the ARM Cortex ecosystem supports.

In the longer term, Xilinx is developing C-to-FPGA compiler flows in an effort to eventually offer software and electronic-system developers a way to readily move functions between software and hardware programming environments to rapidly develop, evaluate and optimize their systems. "The idea is that they will be able to develop in a C-based environment and see results in hardware and software rapidly," said DeHaven. In fact, Xilinx has been diligently monitoring research conducted by the benchmarking and analysis firm BDTI in evaluating the use models of C-level synthesis tools (see related BDTI article in this issue).

While software developers will be able to use any commercial development tools supporting the ARM Cortex-A9, Xilinx plans to bundle its own tools with its new device to help folks get started. Bundled in the tool kit and PCB will be an Eclipse-based integrated development environment, GNU-based compiler, debugger and drivers. "Users will be able to leverage the environment of their choice," said DeHaven. "They will be able to develop with industry tools or Xilinx development tools that support the Cortex-A9 and ARM CoreSight™ debug interfaces."

In addition to ARM native support, Xilinx is also working closely with leading third-party solution providers to develop device-specific software bundles—operating systems and development tools—aimed at engineers using the new device.

Suited for Multiple Vertical Markets

Xilinx created the new architecture at the urging of customers who are looking for a device that is scalable, flexible, upgradeable and will allow them to quickly create derivatives for their needs. The Extensible Processing Platform will allow them to further differentiate their products from competing systems based on fixed-function

ASSPs and ASICs. "We've previewed this to several customers and they are eager to get their hands on the device," said Ratford. "I predict the market reach of this device will be incredible."

For example, Xilinx expects any vertical market that incorporates intelligent video will see immediate benefits from using the new device. Intelligent video requires multiple processing steps such as preprocessing pixel-leveling, which is computationally intensive and thus well served by the parallel-processing capabilities of programmable logic. Intelligent video also requires analytic processing at the element level, which is served using a combination of compatible parallel (programmable-logic) and serial (MPU-based) functions. Meanwhile, application processing at the frame level requires decision, control and communications processing typically performed by MPUs.

Among the specific video markets that stand to gain are automotive driver assistance; consumer multiclass, multifunctional printers; general embedded systems that use scanners; industrial smart cameras, including Internet Protocol surveillance cameras and machine vision, DVRs, medical imaging systems, broadcasting studio cameras and transcoders; and defense-grade night-vision equipment.

One intelligent-video application that will see immediate benefit from the new architecture is automotive driver assistance. Key customers in this space have been urging Xilinx for years to create an ARM-based extensible platform.

Automotive customers will be able to program the device to control and analyze

28-nm FPGAs will target the right mix of power consumption and performance

In February 2010, Xilinx announced it will manufacture its next-generation FPGAs in 28-nanometer high-k metal gate (HKMG) high-performance, low-power process technologies and will use a new, unified ASMBL™ architecture for the devices.

Xilinx evaluated several technologies at the 28-nm node before concluding that the HKMG high-performance, low-power process offered the ideal mix of power consumption and performance for Xilinx's next-generation FPGAs.

For the last four generations of IC processes, static power caused by transistor leakage has become increasingly problematic. At 28 nm, it becomes a first-order effect. Where dynamic power refers to the power a device consumes when it is performing the tasks it was designed to do, static power is wasted juice that increases overall power consumption and produces more heat.

If left unchecked at the 28-nm node, static power draw can account for well over 50 percent of a device's total power draw, drastically taxing power and thermal budgets. A key to the HKMG high-performance, low-power process is the material it uses as a gate dielectric (thermal insulator): instead of silicon dioxide, the technology employs hafnium dioxide. Where a 40-nm silicon dioxide material only afforded a k value (or gate dielectric constant) of 3.9, the 28-nm HKMG high-performance, low-power process will have a k value of 25, which is much more resistant to leakage.

By pairing HKMG with Xilinx's ASMBL (Advanced Silicon Module Block) architecture—which enables Xilinx to rapidly and cost-effectively assemble multiple domain-optimized platforms with an optimal blend of features—Xilinx believes its next-generation devices will have 50 percent lower static power consumption than devices implemented using alternative 28-nm high-performance processes. At the same time, Xilinx is including architectural innovations enabling a system-level performance increase of up to 50 percent over previous-generation FPGAs.

Reducing power consumption also enables increased device capacity. Using this new process technology, Xilinx will introduce a new ultra-high-end 28-nm FPGA to enable applications that require increased capacity in keeping with Moore's Law.

In addition, Xilinx's design team is also introducing extra measures to help designers optimize their designs for the right mix of functionality and performance, as well as power.

Xilinx worked closely with customers to identify and understand the architectural bottlenecks in their systems. One of the biggest areas constraining performance is in interfacing the FPGA with other devices. To address this issue, Xilinx will introduce new clocking technology and will harden critical datapath components. The company will also introduce fine-grained clock-gating features and new place-and-route algorithms in the ISE Design Suite to further reduce power. Fine-grained clock-gating technology is a patented algorithm that analyzes the logic equation and disables wasted logic transitions that do not contribute to the final result. This methodology will help customers effectively reduce the power consumption of their designs by as much as 30 percent.

For more information, visit *http://www.xilinx.com/technology/roadmap/28nm-technology.htm.*

— Mike Santarini

The new device will be processor centric, rather than FPGA centric. 'In a lot of applications you want the controlling software to reprogram the FPGA depending on the application you are running, and there are times where you want the processor to run autonomously from the FPGA fabric.'

the data from multiple sensors positioned 360 degrees around a vehicle, with each sensor performing multiple functions simultaneously. The intelligent control sensor could, for example, have these sensors monitor lanes painted in the road, detect swerving vehicles in adjacent lanes, sync the automobile's speed to that of the vehicle ahead, detect pedestrians and monitor spaces between parked cars to locate adequate parking spots—all simultaneously. A system like this might warn the driver instantly if it detects threats; it could even automatically slow the vehicle down to avoid a collision.

Because such a device would be hardware and software programmable, tier-one vendors could make derivatives of this controller for various car manufacturers and various product lines of each of a carmaker's vehicles—without having to change the overall configuration of the control unit. That capability will save OEMs vast amounts of time, effort and money. Further, the software and hardware programmability means that these devices can be serviced or upgraded in the field.

Similarly, in industrial controls, users can create systems that will manage and analyze data from a series of sensors and motors that in real time can identify defective products flying down an assembly line, detect cracks in the machinery or power down motors when they are running too hot or are not in use to reduce factory costs, optimize operations and even save the lives of workers.

The new device also holds great promise for applications in the wired and wireless communications markets, especially in wireless LTE radio, baseband and enterprise femtocell markets, and in routers, switches and multiplexers in the wired communications arena.

Xilinx anticipates the device will also find multiple uses in the military-aerospace business, especially in cockpit control, munitions and communications equipment supporting the Global Information Grid (see the cover story in *Xcell Journal*, Issue 69).

National Instruments, a longtime customer of Xilinx and alpha customer of the new device, has been closely monitoring the development process and giving feedback to Xilinx. Currently, National Instruments pairs on a PCB a real-time processor with a Xilinx FPGA in its NI Reconfigurable I/O (RIO) LabVIEW FPGA embedded platform (*http://www.ni.com/fpga/*). The platform offers a wide range of peripheral I/O and predefined software libraries that customers can mix and match to create unique embedded systems for various vertical markets. By offloading some of the functions from the standalone processor and instead running them in the FPGA, LabVIEW FPGA can run much faster and with the determinism required for instrumentation, measurement and control applications. The LabVIEW FPGA environment also enables the typical LabVIEW user or market-domain expert to accomplish this without knowing anything about the details of FPGA design.

National Instruments expects embedded products based on the new Xilinx architecture will run much faster, with the added benefit of consuming much less power, said NI R&D fellow Keith Odom.

"Now NI will have a very high-performance processor capable of running our highly productive graphical design environment, and we have very high bandwidth connecting the processor system to the programmable-logic fabric," said Odom. "The amount of data we can transfer between the processor and the programmable-logic fabric is much higher than you would typically see in an FPGA that has an embedded processor or even a microcontroller-based ASSP. Because the bandwidth is so high, you can move beyond mechanical and audio and into more electrical-, radio- or vision-related applications, along with enabling much more sophisticated algorithms to be applied to the data in all applications. This device opens up a lot of new possibilities."

Odom notes that because it essentially integrates two devices in one, data communication on the new device will consume less power. "Because these blocks are connected with so many I/Os, you don't have to burn the power you normally would in high-speed chip-to-chip communication. You'll also be able to power it down into standby modes," he said.

Odom also likes the fact that the new device will be processor centric, rather than FPGA centric. "It is absolutely key," said Odom. "In a lot of applications you want the controlling software to reprogram the FPGA depending on the application you are running, and there are times where you want the processor to run autonomously from the FPGA fabric. So our application constantly switches out what's on the FPGA based on your current processing needs, and this architecture is ideal for that."

"It will be incredible to see what customers will be able to do with device once we've released it," said Ratford of Xilinx. "We're very excited, but still have some work to do for the device to reach its full potential."

Xilinx will announce pricing and availability of the new device in early 2011. To learn more details about the new architecture so you can start development today, visit *www.xilinx.com/technology/roadmap/processing-platform.htm.*